REMARKS

Initially, in accordance with Applicant's duty to provide a summary of an interview, Applicants submit that an interview, with the Examiner, took place on February 18, 2009. Applicants appreciate the courtesies extended by the Examiner during that interview. During the interview, Applicants' representative presented arguments distinguishing the claimed invention over the Dinker et al. (U.S. Patent No. 7,206,836) and Bobbitt et al. (U.S. Patent Application Publication No. 2003/0115218) references. The Examiner agreed to reconsider the rejection in response to the filing of this after final response. No formal agreement was reached.

Further, in the final Office Action, the Examiner provided a section entitled "Claim Clarifications - 35 USC § 101" in which the Examiner made certain assertions regarding Applicants' claims and specification (final Office Action, paragraph 4). Applicants' silence with regard to these assertions should not be construed as Applicants' acquiescence in these assertions.

In the final Office Action, the Examiner rejected claims 1, 3-11, 13, 15, 16, 19-24, 27, 30, 31, and 34-39 under 35 U.S.C. § 103(a) as unpatentable over Dinker et al. in view of Bobbitt et al.; and rejected claims 25, 26, 28, 29, 32, and 33 under 35 U.S.C. § 103(a) as unpatentable over Dinker et al. in view of Bobbitt et al. and Rao et al. (U.S. Patent No. 5,689,706).

Applicants respectfully traverse the Examiner's rejections under 35 U.S.C. § 103. Claims 1, 3-11, 13, 15, 16, and 19-39 remain pending.

REJECTION UNDER 35 U.S.C. § 103 BASED ON DINKER ET AL. AND BOBBITT ET AL.

In paragraph 6 of the Office Action, the Examiner rejected claims 1, 3-11, 13, 15, 16, 19-24, 27, 30, 31, and 34-39 under 35 U.S.C. § 103(a) as allegedly unpatentable over Dinker et al.

in view of Bobbitt et al. Applicants traverse the rejection.

Independent claim 1, for example, is directed to a file system that comprises a plurality of servers configured to store file data as chunks; and a master connected to the servers and configured to store namespace data that includes file identifiers for files for which the file data is stored as chunks, store mapping data that maps the file identifiers to the chunks to which the file identifiers correspond, store an operation log that includes a record of changes to at least one of the namespace data or the mapping data, and store location data that identifies which of the servers stores which of the chunks, where the master is configured to communicate with the servers at startup of the master to identify the chunks stored by the servers, and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data.

Dinker et al. and Bobbitt et al., whether taken alone or in any reasonable combination, do not disclose or suggest the combination of features recited in amended claim 1. For example, Dinker et al. and Bobbitt et al. do not disclose or suggest a master that is configured to, among other things, store an operation log that includes a record of changes to at least one of namespace data, which includes file identifiers for files for which the file data is stored as chunks, or mapping data, which maps the file identifiers to the chunks to which the file identifiers correspond, as recited in claim 1.

The Examiner admitted that Dinker et al. does not disclose or suggest an operation log, but alleged that Bobbitt et al. discloses storing an operation log that includes a record of changes to at least one of namespace data or mapping data, and cited paragraphs 0048 and 0052-0054 of Bobbitt et al. for support (final Office Action, pages 3-4). Applicants submit that the disclosure

of Bobbitt et al. provides no support for the Examiner's allegation.

At paragraph 0048, Bobbitt et al. discloses:

> Configuration information 45 includes configuration data that identifies what physical server(s) the various gtrees for a given GVV are hosted on, what physical devices the master and slave gtrees are stored on, the exports each server provides, and the roles played by the various components in a Gossamer virtual file system. Configuration information also may include schedule data (i.e., data pertaining to when migrations are to be performed or considered, when backups are to occur, when the background consistency checker may run, etc.), status files pertaining to operations in progress, such as migration and backup operations, and log files. The configuration information may be stored on one of the servers used to store the master gtrees and/or the slave gtrees, including file server 20, or may be stored on a separate server that is not used to store file system data files that are part of a GVV.

In this section, Bobbitt et al. discloses configuration information that identifies what server hosts various gtrees, what devices store the master and slave gtrees, the exports provided by each server, the roles the various components play, schedule data, status files pertaining to operations in progress, and log files. While this section of Bobbitt et al. mentions "log files," nowhere does Bobbitt et al. disclose or remotely suggest that these log files include a record of changes to namespace data that includes file identifiers for files for which the file data is stored as chunks. Bobbitt et al. also does not disclose or remotely suggest that these log files include a record of changes to mapping data that maps the file identifiers to the chunks to which the file identifiers correspond. Thus, Bobbitt et al. does not disclose or suggest a master that is configured to, among other things, store an operation log that includes a record of changes to at least one of namespace data, which includes file identifiers for files for which the file data is stored as chunks, or mapping data, which maps the file identifiers to the chunks to which the file identifiers correspond, as recited in claim 1.

At paragraphs 0052-0054, Bobbitt et al. discloses a user-view tree that corresponds to a virtual directory and file hierarchy. Bobbitt et al. discloses that translations between virtual

pathname and actual server-pathname are handled through a GVV master directory structure that logically divides its data into three spaces: a Gossamer namespace, a temporary migrating space, and a garbage space. Nowhere does Bobbitt et al. disclose or suggest anything that can reasonably correspond to an operation log that includes a record of changes to namespace data and/or mapping data. Thus, Bobbitt et al. does not disclose or suggest a master that is configured to, among other things, store an operation log that includes a record of changes to at least one of namespace data, which includes file identifiers for files for which the file data is stored as chunks, or mapping data, which maps the file identifiers to the chunks to which the file identifiers correspond, as recited in claim 1.

In response to these arguments, the Examiner alleged that an example of a namespace is a directory and the gtrees of Bobbitt et al. are considered to represent a directory (final Office Action, page 16). The Examiner further alleged that Bobbitt et al. deals with the migration of files, which means that files are being moved from one server to another server, and when files are migrated, their namespace data and mapping data are both updated at the master (final Office Action, page 16). The Examiner alleged that this updating at the master corresponds to an operation log (final Office Action, page 16). Applicants submit that the Examiner's allegations lack merit.

Even assuming, for the sake of argument, that the gtrees, disclosed by Bobbitt et al., can reasonably be interpreted as including namespace data and mapping data, as alleged by the Examiner (a point that Applicants do not concede), the migration of files would not result in the updating of the namespace data and the mapping data, as alleged by the Examiner. As recited in claim 1, namespace data includes file identifiers for files for which file data is stored as chunks.

Contrary to the Examiner's allegation, migrating a file would have no bearing on the namespace data. The migrated file would continue to exist in the file system and only its location would change. Thus, the namespace data corresponding to the file would remain unchanged. As recited in claim 1, mapping data maps file identifiers to the chunks to which the file identifiers correspond. Contrary to the Examiner's allegation, migrating a file would have no bearing on the mapping data. The migrated file (if stored as chunks) would continue to be mapped to the same chunks, though the location of one or more of the chunks might change. Thus, the mapping data corresponding to the file would remain unchanged. For at least these reasons, Applicants submit that the Examiner's allegations lack merit.

Further, even assuming, for the sake of argument, that the disclosure of Bobbitt et al. could reasonably be interpreted as disclosing updating namespace data and mapping data when a file is migrated, as alleged by the Examiner (a point that Applicants do not concede), Bobbitt et al. still would not disclose an operation log that includes a record of changes to at least one of the namespace data or the mapping data. Rather, under the Examiner's allegation, namespace data and mapping data would be updated when a file is migrated. Simply updating namespace data and mapping data is a very different function with a completely different result from storing a record of changes to at least one of the namespace data or the mapping data. Thus, Bobbitt et al. does not disclose or suggest a master that is configured to, among other things, store an operation log that includes a record of changes to at least one of namespace data, which includes file identifiers for files for which the file data is stored as chunks, or mapping data, which maps the file identifiers to the chunks to which the file identifiers correspond, as recited in claim 1.

Dinker et al. and Bobbitt et al., whether taken alone or in any reasonable combination,

also do not disclose or suggest a master that is configured to, among other things, communicate with the servers at startup of the master to identify the chunks stored by the servers and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as further recited in claim 1.

The Examiner alleged that Dinker et al. discloses these features and cited column 6, lines 8-67, of Dinker et al. for support (final Office Action, page 3). Applicants submit that the disclosure of Dinker et al. provides no support for the Examiner's allegation.

At column 6, lines 8-67, Dinker et al. discloses a replication topology manager that maintains the distribution of data on the nodes, as defined by a replication topology. Dinker et al. discloses that the replication topology manager can initiate one or more copy operations by the nodes so that the replication of data within the cluster conforms to the replication topology. Even assuming, for the sake of argument, that the nodes disclosed by Dinker et al. correspond to servers and that the replication topology manager corresponds to a master (points that Applicants do not concede), nowhere in this section, or elsewhere, does Dinker et al. disclose or suggest that the replication topology manager communicates with the nodes at startup to identify the data stored by the nodes. In fact, Dinker et al. does not specifically disclose that the replication topology manager performs any communication with the nodes to identify the data stored by the nodes, whether at startup or at another time. Thus, Dinker et al. does not disclose or suggest a master that is configured to, among other things, communicate with the servers at startup of the master to identify the chunks stored by the servers and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as recited in claim 1.

Further, nowhere in the above-identified section, or elsewhere, does Dinker et al. disclose or suggest that the replication topology manager records, in a non-persistent manner, information regarding data stored by each of the nodes. Thus, Dinker et al. does not disclose or suggest a master that is configured to, among other things, communicate with the servers at startup of the master to identify the chunks stored by the servers and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as recited in claim 1.

The Examiner alleged that the communication interface, in Dinker et al., notifies the replication topology manager whenever changes in cluster membership are detected and, therefore, it is inherent that when a node, which is to become a manager, is added to the cluster, the node will have to receive the topology information (final Office Action, page 17). Applicants submit that the Examiner's allegation lacks merit and falls short of establishing a rejection based on inherency.

According to M.P.E.P. § 2112, the fact that a certain result or characteristic may occur or be present in the prior art is not sufficient to establish the inherency of that result or characteristic. In relying upon the theory of inherency, the Examiner must provide a basis in fact and/or technical reasoning to reasonably support the determination that the allegedly inherent characteristic necessarily flows from the teachings of the applied prior art. Inherency cannot be established by probabilities or possibilities. The mere fact that a certain thing may result from a given set of circumstances is not sufficient. In this case, the Examiner's allegation does not meet the requisite burden of proof to establish inherency. For example, the Examiner has provided absolutely no evidence that when a new node is added to a cluster in Dinker et al., that new node

would necessarily communicate with the other nodes of the cluster, at startup of the new node, to identify the data stored by the nodes and record, in a non-persistent manner, information regarding the data stored by each of the nodes as the location data. Thus, contrary to the Examiner's allegation, Dinker et al. does not disclose or suggest a master that is configured to, among other things, communicate with the servers at startup of the master to identify the chunks stored by the servers and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as recited in claim 1.

The disclosure of Bobbitt et al. does not cure the deficiencies in the disclosure of Dinker et al. For example, Bobbitt et al. does not disclose or suggest a master that is configured to, among other things, communicate with the servers at startup of the master to identify the chunks stored by the servers and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as recited in claim 1. Bobbitt et al. appears to actually teach away from these features by disclosing the persistence of files when the computer turns on and off (paragraph 0027).

For at least these reasons, Applicants submit that claim 1 is patentable over Dinker et al. and Bobbitt et al., whether taken alone or in any reasonable combination. Claims 3-11 and 19-24 depend from claim 1 and are, therefore, patentable over Dinker et al. and Bobbitt et al. for at least the reasons given with regard to claim 1. Claims 3-11 and 19-24 are also patentable for reasons of their own.

For example, claim 11 recites information that includes version numbers of the chunks. Dinker et al. and Bobbitt et al. do not disclose or suggest this feature.

The Examiner alleged that Dinker et al. discloses this feature and cited column 6, lines 8-

67, of Dinker et al. for support (final Office Action, page 6). Applicants submit that the disclosure of Dinker et al. provides no support for the Examiner's allegation.

At column 6, lines 8-67, Dinker et al. discloses a replication topology manager that maintains the distribution of data on the nodes, as defined by a replication topology. Dinker et al. discloses that the replication topology manager can initiate one or more copy operations by the nodes so that the replication of data within the cluster conforms to the replication topology. Dinker et al. does not disclose that versions of the data are maintained. Thus, Dinker et al. does not disclose or suggest information that includes version numbers of the chunks, as recited in claim 11.

For at least these additional reasons, Applicants submit that claim 11 is patentable over Dinker et al. and Bobbitt et al.

Dependent claim 22 recites that the chunk handle encodes a timestamp. Dinker et al. and Bobbitt et al. do not disclose or suggest this feature.

The Examiner alleged that Bobbitt et al. discloses this feature and cited paragraph 0054 of Bobbitt et al. for support (final Office Action, page 10). Applicants submit that the disclosure of Bobbitt et al. provides no support for the Examiner's allegation.

At paragraph 0054, Bobbitt et al. discloses a globally unique identifier (GUID) assigned to a file. Nowhere does Bobbitt et al. disclose or remotely suggest that the GUID encodes a timestamp. Rather Bobbitt et al. discloses that the GUID is merely a 128-bit identifier (paragraph 0055). Thus, Bobbitt et al. does not disclose or suggest a chunk handle that encodes a timestamp, as recited in claim 22.

For at least these additional reasons, Applicants submit that claim 22 is patentable over

Dinker et al. and Bobbitt et al.

Dependent claim 23 recites that the master is configured to update the location data by periodically instructing the servers to provide information regarding the chunks stored by the servers. Dinker et al. and Bobbitt et al. do not disclose or suggest this feature.

The Examiner alleged that Dinker et al. discloses this feature and cited column 6, lines 8-67, of Dinker et al. for support (final Office Action, page 10). Applicants submit that the disclosure of Dinker et al. provides no support for the Examiner's allegation.

At column 6, lines 8-67, Dinker et al. discloses a replication topology manager that maintains the distribution of data on the nodes, as defined by a replication topology. Dinker et al. discloses that the replication topology manager can initiate one or more copy operations by the nodes so that the replication of data within the cluster conforms to the replication topology. Even assuming, for the sake of argument, that the nodes disclosed by Dinker et al. correspond to servers and that the replication topology manager corresponds to a master (points that Applicants do not concede), nowhere in this section, or elsewhere, does Dinker et al. disclose or suggest that the replication topology manager periodically instructs the nodes to provide information regarding the data stored by the nodes. In fact, Dinker et al. does not specifically disclose that the replication topology manager performs any communication with the nodes to identify the data stored by the nodes. Thus, Dinker et al. does not disclose or suggest a master that is configured to update the location data by periodically instructing the servers to provide information regarding the chunks stored by the servers, as recited in claim 23.

For at least these additional reasons, Applicants submit that claim 23 is patentable over Dinker et al. and Bobbitt et al.

Dependent claim 24 recites that the operation log includes a logical timeline that defines an order for concurrent operations. Dinker et al. and Bobbitt et al. do not disclose or suggest this feature.

The Examiner alleged that Bobbitt et al. discloses this feature and cited paragraph 0048 of Bobbitt et al. for support (final Office Action, page 11). Applicants submit that the disclosure of Bobbitt et al. provides no support for the Examiner's allegation.

At paragraph 0048, Bobbitt et al. discloses configuration information that identifies what server hosts various gtrees, what devices store the master and slave gtrees, the exports provided by each server, the roles the various components play, schedule data, status files pertaining to operations in progress, and log files. As explained above with regard to claim 1, while this section of Bobbitt et al. mentions "log files," nowhere does Bobbitt et al. disclose or remotely suggest that these log files include a record of changes to namespace data, which includes file identifiers for files for which the file data is stored as chunks, and/or mapping data, which maps the file identifiers to the chunks to which the file identifiers correspond. Thus, Bobbitt et al. cannot disclose or suggest an operation log includes a logical timeline that defines an order for concurrent operations, as recited in claim 24.

For at least these additional reasons, Applicants submit that claim 24 is patentable over Dinker et al. and Bobbitt et al.

Independent claims 13, 15, 16, and 30 recite features similar to, yet possibly different in scope from, features described above with regard to claim 1. Claims 13, 15, 16, and 30 are, therefore, patentable over Dinker et al. and Bobbitt et al., whether taken alone or in any reasonable combination, for at least reasons similar to reasons given with regard to claim 1.

Claims 27 and 38 depend from claim 13 and are, therefore, patentable over Dinker et al. and Bobbitt et al. for at least the reasons given with regard to claim 13. Claims 34 and 35 depend from claim 15 and are, therefore, patentable over Dinker et al. and Bobbitt et al. for at least the reasons given with regard to claim 15. Claims 36 and 37 depend from claim 16 and are, therefore, patentable over Dinker et al. and Bobbitt et al. for at least the reasons given with regard to claim 16. Claims 31 and 39 depend from claim 30 and are, therefore, patentable over Dinker et al. and Bobbitt et al. for at least the reasons given with regard to claim 31. Claims 27 and 31 also recite a feature similar to a feature recited in claim 24. Therefore, claims 27 and 31 are also patentable over Dinker et al. and Bobbitt et al. for at least reasons similar to the reasons given with regard to claim 24.

Accordingly, Applicants respectfully request reconsideration and withdrawal of the rejection of claims 1, 3-11, 13, 15, 16, 19-24, 27, 30, 31, and 34-39 under 35 U.S.C. § 103 based on Dinker et al. and Bobbitt et al.

REJECTION UNDER 35 U.S.C. § 103 BASED ON DINKER ET AL., BOBBITT ET AL., AND RAO ET AL.

In paragraph 7 of the final Office Action, the Examiner rejected claims 25, 26, 28, 29, 32, and 33 under 35 U.S.C. § 103(a) as allegedly unpatentable over Dinker et al. in view of Bobbitt et al. and Rao et al. Applicants traverse the rejection.

Claims 25 and 26 depend from claim 1, claims 28 and 29 depend from claim 13, and claims 32 and 33 depend from claim 30. Without acquiescing in the Examiner's rejection with regard to claims 25, 26, 28, 29, 32, and 33, Applicants respectfully submit that the disclosure of Rao et al. does not cure the deficiencies in the disclosures of Dinker et al. and Bobbitt et al.

identified above with regard to claims 1, 13, and 30. For example, Rao et al. does not disclose or suggest an operation log that includes a record of changes to at least one of the namespace data and/or the mapping data, as recited in claims 1 and 13; or an operation log that includes a record of changes to the namespace data and the mapping data, as recited in claim 30. Rao et al. also does not disclose or suggest a master that is configured to communicate with the servers to identify the chunks stored by the servers, and record, in a non-persistent manner, information regarding the chunks stored by each of the servers as the location data, as recited in claim 1; means for communicating with the servers to identify file data stored by the servers as chunks and means for storing, in a non-persistent manner, location information that identifies ones of the servers that store the chunks, as recited in claim 13; or communicating with the server devices to identify file data stored by the server devices as chunks and storing location information that identifies ones of the server devices that store the chunks, as recited in claim 30.

Therefore, claims 25, 26, 28, 29, 32, and 33 are patentable over Dinker et al., Bobbitt et al., and Rao et al., whether taken alone or in any reasonable combination, for at least the reasons given with regard to claims 1, 13, and 30.

Accordingly, Applicants respectfully request reconsideration and withdrawal of the rejection of claims 25, 26, 28, 29, 32, and 33 under 35 U.S.C. § 103 based on Dinker et al., Bobbitt et al., and Rao et al.

CONCLUSION

In view of the foregoing amendments and remarks, Applicants respectfully request the Examiner's reconsideration of the application and the timely allowance of pending claims 1, 3-11, 13, 15, 16, and 19-39.

As Applicants' remarks with respect to the Examiner's rejections are sufficient to overcome these rejections, Applicants' silence as to certain assertions by the Examiner in the Office Action or certain requirements that may be applicable to such rejections (e.g., whether a reference constitutes prior art, reasons for modifying a reference and/or combining references, assertions regarding dependent claims, etc.) is not a concession by Applicants that such assertions are accurate or such requirements have been met, and Applicants reserve the right to analyze and dispute these assertions/requirements in the future.

If the Examiner does not believe that all pending claims are now in condition for allowance, the Examiner is urged to contact the undersigned to expedite prosecution of this application.

To the extent necessary, a petition for an extension of time under 37 C.F.R. § 1.136 is hereby made. Please charge any shortage in fees due in connection with the filing of this paper, including extension of time fees, to Deposit Account No. 50-1070 and please credit any excess fees to such deposit account.

Respectfully submitted,

HARRITY & HARRITY, LLP

By: /Paul A. Harrity, Reg. No. 39574/
Paul A. Harrity
Reg. No. 39,574

Date: February 20, 2009
11350 Random Hills Road
Suite 600
Fairfax, Virginia 22030
(571) 432-0800
Customer No. 44989